arXiv:1509.08124v2 [stat.ME] 25 Feb 2016

A testing-based approach to the discovery of differentially correlated variable sets.

Kelly Bodwin, Kai Zhang, and Andrew Nobel

February 26, 2016

Abstract

Given data obtained under two sampling conditions, it is often of interest to identify variables that behave differently in one condition than in the other. We introduce a method for differential analysis of second-order behavior called Differential Correlation Mining (DCM). The DCM method identifies differentially correlated sets of variables, with the property that the average pairwise correlation between variables in a set is higher under one sample condition than the other. DCM is based on an iterative search procedure that adaptively updates the size and elements of a candidate variable set. Updates are performed via hypothesis testing of individual variables, based on the asymptotic distribution of their average differential correlation. We investigate the performance of DCM by applying it to simulated data as well as recent experimental datasets in genomics and brain imaging.

Keywords: differential correlation mining, association mining, biostatistics, genomics, high dimensional data

Address for correspondence: Kelly Bodwin, *kbodwin@email.unc.edu*

1 Introduction

In many statistical problems one has two datasets that measure the same variables under different conditions. In the analysis of such data, it is common to assume that the samples in each dataset are generated from two underlying distributions. Even when the data is high dimensional, differences between the distributions may only be present for a small number of variables, and it is often of interest to identify these key variables. In this paper, we present a new method of second order comparative analysis, called Differential Correlation Mining (DCM), which identifies sets of variables such that the average pairwise correlation between variables in the set is higher in one sample condition than in another. The method does not make use of auxiliary information, apart from the separation of samples into pre-determined groups (e.g. treatment vs control). DCM is applicable to both low and high dimensional datasets.

Most often, differential behavior between sample groups is measured by *first-order* statistics, which are functions of a single variable. Familiar first-order statistics include the sample mean and the sample variance. A well-studied example of first-order differential analysis is the study of differential gene expression in microarrays (see Cui and Churchill (2003) for a canonical example, or Soneson and Delorenzi (2013) and the references therein for an overview of several methods). Other applications of first-order differential analysis include text analysis for authorship identification (Stamatatos, 2009), studies of brain functionality based on regional activation (Phan et al., 2002), and investigation of cultural bias in standardized testing (Wainer and Braun, 2013).

The use of first-order statistics only allows for analysis of a single variable at a time. To study relationships between pairs of variables, one requires functions of two variables, which specify *second-order* statistics. Examples of second-order statistics include correlation, covariance, and distance. When one wishes to understand interactions between many variables (as in clustering problems), data may be summarized in matrix form, where each entry in the matrix represents the observed value of a second-order statistic. It is common to look within a matrix of relational data for groups of variables that have high pairwise association. Interconnected variable sets may represent, e.g., social groups in a communication network (Lewis et al., 2008), genes in common protein pathways (Jiang et al., 2004), functionally similar brain regions (Greicius et al., 2002).

While there is a large literature on clustering and networks, to the best of our knowledge, there is relatively little work comparing second-order behavior across two sample conditions. The many insights obtained from ordinary second-order variable set selection lead us to believe that a second-order differential approach will be of scientific interest. The methods introduced in this paper fall under the broader heading of *differential association mining*. As in ordinary association mining, we are interested in the pairwise behavior of variables; however, in the differential setting, we must consider two different relational matrices. In some cases, simply taking the difference of the matrices and applying ordinary clustering methods would suffice. However, most second-order statistics - including

the focus of this paper, the linear correlation coefficient - require a more careful treatment. For instance, two sample correlation matrices will exhibit vastly different random behavior based on the sample sizes of the corresponding datasets, and will have a complex dependency structure when the population correlation matrix is not the identity.

The DCM method proposed here addresses differential correlation mining in a direct way. (Section 1.2 considers possible alternatives based on existing work.) DCM seeks variable sets that form differentially correlated (DC) cliques. In a graph, a clique is a set of nodes that is fully connected, in the sense that there is an edge between every pair of nodes in the set. Informally, a DC clique is a set of variables such that each variable in the set has a positive (usually large) average differential correlation with the other variables in the set.

More formally, let $\mathbf{R}_1, \mathbf{R}_2$ be the $p \times p$ population correlation matrices of the distributions underlying sampling conditions 1 and 2, respectively. Let $J \subset [p]$, where $[p]$ is the index set $\{1, ..., p\}$, and define

$$\Delta(i, J) \;=\; \frac{1}{|J|} \sum_{j \in J} (\mathbf{R}_1 - \mathbf{R}_2)_{ij} \tag{1}$$

to be the difference of correlations between variable i and variables in index set J. Here the subscript ij denotes the element in the i-th row and j-th column of the corresponding matrix, and $|J|$ is the cardinality of the set J. We define DC cliques as follows.

Definition 1.1. *Let $\mathbf{R}_1, \mathbf{R}_2$ be given and let $\Delta(\cdot, \cdot)$ be defined as in (1). An index set $A \subseteq [p]$ with at least two elements is a* DC clique *for $\mathbf{R}_1 - \mathbf{R}_2$ if*

1. *$\Delta(i, A) > 0$ iff $i \in A$,*
2. *The set* A *cannot be written as a disjoint union of nonempty index sets $A_1, A_2 \subset [p]$ such that A_1 and A_2 satisfy condition 1 above.*

Condition 1 ensures that no relevant variables are omitted from a DC clique (every variable that is positively differentially correlated relative to the set A is included in A) and that a DC clique does not contain any extraneous elements. Condition 1 implies that a DC clique has larger average pairwise correlation under the first distribution than under the second. Condition 2 ensures that a DC clique cannot be subdivided into two smaller DC cliques. Importantly, the definition places *no conditions* on the correlation matrices $\mathbf{R}_1$ and $\mathbf{R}_2$. In particular, $\mathbf{R}_1$ and $\mathbf{R}_2$ need not be sparse, and need not satisfy any structural constraints such as bandedness. For a given pair $\mathbf{R}_1, \mathbf{R}_2$, it may happen that no DC cliques exist, or that the entire variable set forms a DC clique. Note that the definition of DC cliques is not symmetric: in general, the DC cliques for $\mathbf{R}_1 - \mathbf{R}_2$ will be different from those for $\mathbf{R}_2 - \mathbf{R}_1$.

As defined above, DC cliques are features of the underlying population distributions of the data. The broad objective of DCM is to use observed data to identify DC cliques,

or approximations of these, without prior knowledge of the identity, number, or size of the DC cliques present in the population. It is worth noting that the DCM algorithm and supporting analysis described here is easily adapted to a non-differential correlation mining algorithm. Code for the correlation mining procedure is available in the DCM packages for R and Matlab.

Remark. Some bioinformatics literature uses the phrase "Differential Co-Expression", sometimes abbreviated "DC", as an umbrella term for all differential second-order gene expression behavior. In this paper, "DC" will refer specifically to differential correlation; when a distinction must be made with co-expression or covariance, this will be made explicit.

1.1 An example

To motivate our definition of DC cliques, we provide an illustrative real-world example. Figure 1 shows an empirical DC clique identified by DCM in real data from The Cancer Genome Atlas (TCGA) Research Network (http://cancergenome.nih.gov/). The two sample conditions under consideration are Her-2 type breast cancer tumors and Luminal B type tumors, as classified by Perou et al. (2000). (Further results for the TCGA dataset are provided in Section 5.)

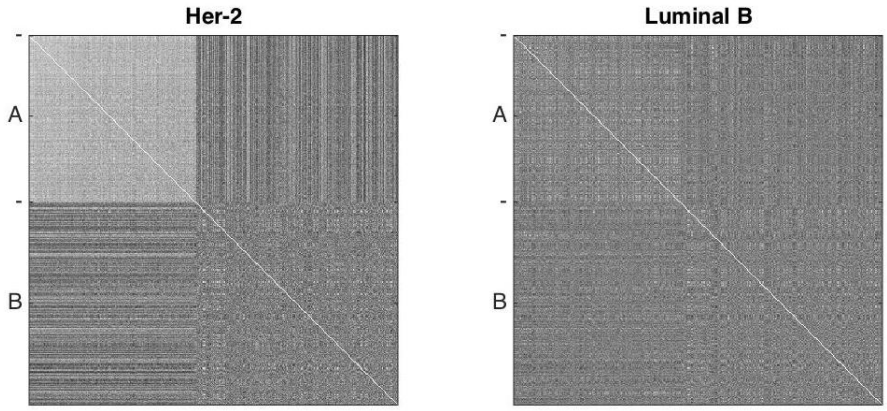

Figure 1: Sample correlation matrices for each of two breast cancer tumor subtypes, showing observed DC clique (A) and random genes (B) .

Figure 1 shows the sample correlation matrices within each tumor type, restricted to a set of 365 variables consisting of an empirical DC clique of size 165 selected by DCM (A), and 200 randomly chosen variables (B). The variables B are included for contrast, and to show that the differential correlation observed in A is not present in the entire dataset.

The figure illustrates the second-order behavior and the differential nature of the identified DC clique A. The block pattern in the upper left corner of the Her-2 matrix shows that every entry in the correlation matrix of A is large, suggesting that all the variables of A are strongly pairwise correlated. The Luminal B sample correlation shows a similar pattern, but it is much less pronounced. No such pattern is seen among the variables in B.

In general, the results of DCM are distinct from those found by first-order analysis (e.g. differential expression). For example, Figure 2 shows the relative differential expression, overall expression level, and differential variation for the above estimated DC clique A. For this plot, we ranked all genes in the study ($p = 15,785$) by (a) t statistic of differential mean expression between Her-2 and Luminal B samples, (b) overall expression in Her-2 samples, and (c) ratio of sample variations (F statistic) for Her-2 versus Luminal B samples. The histograms in Figure 2 show the ranking of the genes in A. The overall uniformity of the histograms indicates that the variables in the observed DC clique A do *not* exhibit standard first-order differential behavior. Similar results were observed for all other data studied in this paper.

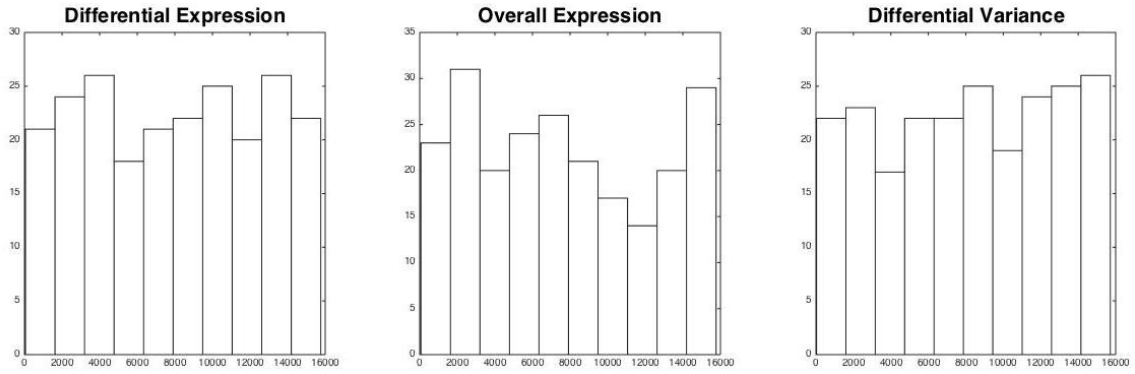

Figure 2: Ranks of genes in observed DC clique (A) out of 15,785 total genes.
(Ranked by: Differential expression, as measured by p-values of 2-sample t-tests; mean overall expression among Her-2 samples; and difference of sample variance between Her-2 and Luminal B.)

By targeting DC cliques, the DCM method identifies variables whose *joint* behavior is different across sample conditions. The results are readily interpretable as sets of variables that interact strongly under one sample condition but only weakly (or not at all) under another. In this paper, we will demonstrate DCM is an effective and efficient way to identify differentially correlated variable sets from observed data.

1.2 Related work

There is a substantial body of recent work concerning estimation and testing of covariance and correlation matrices in high-dimensional settings. However, to the best of our knowl-

edge, none of this work addresses the search for DC cliques or related differential structures in a direct way. Below we provide an overview of work that is related to DCM. In what follows, let $\mathbf{R}_1, \mathbf{R}_2$ denote the population correlation matrices of two data distributions, and let $\widehat{\mathbf{R}}_1, \widehat{\mathbf{R}}_2$ denote the corresponding sample correlation matrices.

Mining from single correlation matrices. Non-differential correlation mining, in which one searches for highly associated variables from a single dataset, has been well-studied, typically in the context of clustering. Kriegel et al. (2009) provides a survey of clustering methods for high-dimensional data based on correlation distance. Datta and Datta (2002) and Jiang et al. (2004) and the references therein give an overview of methods developed specifically for clustering of gene expression. In general, typical clustering or community detection methods must be adapted for application to correlation distances to correct for bias (see e.g. MacMahon and Garlaschelli (2015) for an illustrative example).

Estimation and hypothesis testing. There has been a great deal theoretical work devoted to testing equality of high-dimensional covariance and correlation matrices. When the sample size n is substantially larger than the dimension p, classical results are applicable, e.g., likelihood ratio tests as discussed in Anderson (1959) and Muirhead (1982), or results like those of (Steiger, 1980) for testing individual sample correlation. In the high-dimensional ($p > n$) setting, Cai et al. (2010); Cai and Jiang (2011); Cai et al. (2014) have developed minimax rate optimal tests for the equality of covariance matrices under sparsity assumptions. Results for correlation (rather than covariance) are less prevalent; recent work includes tests for sets of sample correlation coefficients (Donner and Zou, 2014), tests for rank-based correlation matrices (Zhou et al., 2015), and tests for detecting overall dependence (Bassi and Hero, 2012).

In some cases, optimal testing procedures can inform methods for estimation of high-dimensional covariance and correlation matrices. Particularly relevant is the work of Cai and Zhang (2014), which yields an estimator for the difference matrix $\mathbf{D} = \mathbf{R}_1 - \mathbf{R}_2$. This estimator is implemented and discussed further in Section 4. Other approaches to high-dimensional estimation include the work of Bickel and Levina (2008), who discuss a thresholding estimator for covariance matrices; Peng et al. (2008), who estimate partial correlations in sparse regression models; and Rajaratnam et al. (2008), who make use of graphical model techniques for covariance matrix estimation.

Detection of isolated changes in correlation structure. Existing approaches to differential correlation mining are based on examining individual variables for changes in second-order structure across two sample conditions. For example, one may treat $\widehat{\mathbf{R}}_1$ and $\widehat{\mathbf{R}}_2$ as the adjacency matrices of two fully connected, weighted networks, and then look for variables whose connectivity pattern is very different across the two networks (Xia et al., 2014; Gill et al., 2010). Most methods approach differential correlation mining by developing a statistic to measure the change in pairwise correlations of an individual variable. Hu et al. (2010) uses the covariance distance (total difference of covariances); Choi and Kendziorski (2009) uses a direct difference of sample correlations; Fukushima (2013) uses the difference of Fisher transformed sample correlations; and Liu et al. (2010) use a filtration (or

thresholding) step before summing square correlation differences. These methods then permute samples across the two classes to measure the significance of the original differential correlation. Significant variables may then be selected by an appropriate multiple testing procedure.

1.3 Outline

This remainder of the paper is organized as follows. In the next section, we describe in detail the three main steps of the DCM procedure. Section 3 provides a closer examination of the test statistic used in the procedure, including a discussion of its asymptotic distribution. We apply DCM to simulated data in Section 4, and compare the results to possible alternative procedures based upon existing work. Finally, in Section 5 we present the results of two applications of DCM, to the aforementioned TCGA dataset and to brain activity data from the multi-institutional Human Connectome Project.

2 The DCM Procedure

In this section, we present details of the three components of the proposed DCM procedure: initialization, set update, and residualization. The initialization step employs a simple greedy algorithm to select an initial variable set A. Once the initial set is determined, it is passed to an update algorithm that iteratively refines the set, making use of a hypothesis testing framework to test variables for differential correlation. When an estimated DC clique is found, the residualization step prepares the data for further search by removing the differential correlation of the discovered set.

The DCM procedure is summarized below. For detailed pseudocode, see Appendix B.

THE DCM PROCEDURE

- ▷ *Initialization:* Identify a good initial variable set A using a greedy algorithm that identifies a local maximum of a simple score function.
- ▷ *Iteration:* Refine the initial set A. At each iterative step, repeat the following until termination.
  - ▷ *Test* the differential correlation of each variable i with respect to A. Let A' be the set of variables with significant differential correlation, as determined by and FDR controlling multiple testing procedure.
  - ▷ *Terminate* if $A' = A$ or a cycle is observed.
  - ▷ *Update:* Set A to be A'.
- ▷ *Return:* Output variable set A.
- ▷ *Residualization:* Remove the effect of the DC clique A and repeat search with new initial set.

We now provide a more in-depth discussion of each step of the procedure.

2.1 Notation

We assume that the data under condition 1 consists of n_1 independent samples drawn from a distribution F_1 with correlation matrix $\mathbf{R}_1$, and that the data under condition 2 consists of n_2 independent samples drawn from a distribution F_2 with correlation matrix $\mathbf{R}_2$. Let $\mathbb{X}_1 = (\mathbf{U}_1, ..., \mathbf{U}_p) \in \mathbb{R}^{n_1 \times p}$ and $\mathbb{X}_2 = (\mathbf{V}_1, ..., \mathbf{V}_p) \in \mathbb{R}^{n_2 \times p}$ denote the resulting data matrices in standard sample-by-variable form. Thus $\mathbf{U}_j \in \mathbb{R}^{n_1}$ denotes the measurements of variable j under condition 1, while $\mathbf{V}_j \in \mathbb{R}^{n_2}$ denotes the measurements of variable j under condition 2. Let $\mathbb{X}_{1,A} = (\mathbf{U}_j)_{j \in A}$ and $\mathbb{X}_{2,A} = (\mathbf{V}_j)_{j \in A}$ denote the restriction of $\mathbb{X}_1$ and $\mathbb{X}_2$, respectively, to a variable set $A \subset [p]$. Similarly, let $\mathbf{R}_{1,A}$ and $\mathbf{R}_{2,A}$ denote the correlation matrices of the restricted datasets.

Let $\tilde{\mathbf{U}}_j$ and $\tilde{\mathbf{V}}_j$ be the standardized versions of $\mathbf{U}_j$ and $\mathbf{V}_j$, respectively, and define $\tilde{\mathbb{X}}_1 = (\tilde{\mathbf{U}}_1, ..., \tilde{\mathbf{U}}_p)$ and $\tilde{\mathbb{X}}_2 = (\tilde{\mathbf{V}}_1, ..., \tilde{\mathbf{V}}_p)$. Finally, let $\widehat{\mathbf{R}}_1$ and $\widehat{\mathbf{R}}_2$ denote the usual sample correlation matrices of $\mathbb{X}_1$ and $\mathbb{X}_2$, respectively (and $\widehat{\mathbf{R}}_{1,A}$ and $\widehat{\mathbf{R}}_{2,A}$ those of the appropriate restricted datasets). Thus $\left(\widehat{\mathbf{R}}_1\right)_{ij} = \widehat{\mathrm{cor}}_n\left(\mathbf{U}_i, \mathbf{U}_j\right) = \left(\tilde{\mathbb{X}}_1^T \tilde{\mathbb{X}}_1\right)_{ij}$ and a similar relation holds for $\widehat{\mathbf{R}}_2$.

2.2 Initialization

The set update procedure in the second step of DCM readily identifies variables that are significantly differentially correlated relative to a given variable set A, and is most effective when the initial set of variables exhibits at least low levels of differential correlation. (When applied to a randomly chosen set of variables, the set update procedure typically returns an empty set.) The core search procedure could be run exhaustively, beginning with every variable set $A \subset [p]$, but this is not computationally feasible for data sets of high or moderate dimension. As an alternative, we identify initial variable sets exhibiting a moderate degree of differential expression using a greedy search procedure. We then can pass this initial skeleton clique to the set update process to be fleshed out into a final estimated DC clique.

The initialization procedure seeks a local maximum of the score function

$$S(A) = \sum_{i,j \in A} \left\{ (n_1 - 3)^{1/2} \, \varphi\left(\widehat{\mathbf{R}}_1\right) - (n_2 - 3)^{1/2} \, \varphi\left(\widehat{\mathbf{R}}_2\right) \right\}_{ij}$$

where φ is the element-wise Fisher transformation of sample correlations, namely

$$\varphi(r) = \frac{1}{2} \log\left(\frac{1-r}{1+r}\right).$$

The Fisher transformation and subsequent weighting by degrees of freedom ensures that the first and second terms in the sum have approximately equal variances for each i, j. In this way, we account for the natural variance of the individual elements of $\widehat{\mathbf{R}}_1$ and $\widehat{\mathbf{R}}_2$ as well as possible imbalance in sample sizes across the two datasets.

To find a local maximizer of $S(\cdot)$, we begin with a random seed A. We consider only pairwise swaps, replacing an element of A with one from A^c; the set A is then updated by making the swap that produced the largest increase in the score. Since exactly one element is added and removed at each stage, the size of the variable set remains constant. The cardinality of A is user-specified (with a default of 50). Due to the subsequent set update procedure, we find that a wide range of initial cardinalities result in the same final outcome.[1] Because of the random seeding, the algorithm is not purely deterministic; however, in practice the same local maximum is reached from most seeds.

Further discussion of the initializing algorithm is available as supplemental material; pseudocode for its implementation may be found in Appendix B. A closely related method is implemented in Section 4 for comparison with DCM.

[1] As a rule, erring on the side of initial cardinalities that are *smaller* than the expected output set size is advisable, to avoid drowning out signal with too much noise.

2.3 Core set update procedure

The heart of the DCM procedure is the set update algorithm, which makes use of multiple testing principles to iteratively refine a variable set A. Recall that the goal of DCM is to estimate DC cliques from the data. To this end, the set update procedure is designed to identify variable sets exhibiting the properties of a true DC clique up to a level of statistical significance.

Consider a single iterative step, at which we update a given variable set A. We wish to determine whether each variable i (including those in A itself) ought to be included in the updated set A'. Since our eventual goal is to discover a DC clique, we perform hypothesis tests based upon the principles of Definition 1.1. For fixed A, the tests for variable i may be written as

$$H_0(i, A) : \Delta(i, A) = 0 \qquad \text{vs.} \qquad H_a(i, A) : \Delta(i, A) > 0\,. \tag{2}$$

Recall that $\Delta(i, A)$, as defined in (1), is a difference of average pairwise correlations between i and elements of A from $\mathbf{R}_1$ and $\mathbf{R}_2$, so this amounts to a test of differential correlation *relative* to a fixed set A. Our updated set is then given by

$$A' = \{i \,:\, H_0(i, A) \text{ was rejected by simultaneous multiple hypothesis testing}\}.$$

In order to test the hypotheses in (2), we require a test statistic. A natural choice is the corresponding sample quantity,

$$\hat{\Delta}(i, A) = \frac{1}{|A|} \sum_{j \in A} (\widehat{\mathbf{R}}_1 - \widehat{\mathbf{R}}_2)_{ij}. \tag{3}$$

In addition to being a straightforward choice, this test statistic exhibits several desirable properties discussed in Section 3.

Let $\delta(i, A)$ denote the realized value of the test statistic for a particular dataset. It is clear that large positive values of $\delta(i, A)$ provide support for the alternate hypothesis in (2), while values that are negative or close to zero provide evidence in favor of the null. Thus, in order to test the hypotheses, we calculate a p-value of the form

$$\mathrm{p}(i : A) = \mathbb{P}_0 \left(\hat{\Delta}(i, A) > \delta(i, A) \right) ,$$

where the probability $\mathbb{P}_0$ is the (unknown) distribution of $\hat{\Delta}(i, A)$ under the null hypothesis $\Delta(i, A) = 0$. Since we make no assumptions about the distributions of data (F_1 and F_2), we instead make use of asymptotic results to approximate the above probability. We show in Section 3.2 that, under appropriate regularity assumptions, and for large enough sample sizes n_1, n_2,

$$\mathrm{p}(i : A) \approx 1 - \Phi \left(\frac{\delta(i, A)}{\hat{\sigma}_0} \right) , \tag{4}$$

where $\hat{\sigma}_0^2$ is an estimate of the variance of $\hat{\Delta}(i, A)$.

The collection of p-values $\{\mathrm{p}(i : A)\}_{i=1}^p$ measure the significance of the differential correlation of each variable relative to A. In order to select a set of significant variables A', we apply the modified FDR procedure of Benjamini and Yekutieli to the p-values. Specifically, we carry out the following steps

1. Order the p-values $\{\mathrm{p}(i : A)\}_{i=1}^p$ as $\{\mathrm{p}_{(1)}, ..., \mathrm{p}_{(p)}\}$.

2. Define the cutoff index k^* by

$$k^* = \max \left\{ k \, : \, \mathrm{p}_{(k)} < \left(\sum_{i=1}^{k} 1/i \right)^{-1} \left(\frac{k\alpha}{p} \right) \right\}.$$

3. Let $A' = \{i \, : \, \mathrm{p}(i : A) \leq \mathrm{p}_{(k^*)}\}$.

Recall that we impose no assumptions on the structure of correlation matrices $\mathbf{R}_1$ and $\mathbf{R}_2$. In particular, it is possible that p-values $\mathrm{p}(i : A)$ and $\mathrm{p}(j : A)$ may be negatively correlated, and negative correlation is common in real data sets. Most common multiple testing methods assume independence or positive dependency between p-values. The possibility of negative dependency of p-values necessitates the more conservative multiple testing method of Benjamini and Yekutieli (2001), which controls the expected False Discovery Rate at level α.

The main search procedure terminates when it degenerates ($A = \emptyset$) or converges ($A = A' \neq \emptyset$). For the degenerate case, the interpretation is simple: the initial set (chosen in the first step of the DCM procedure) was not sufficiently differentially correlated to yield an interesting result. In the second case, we have identified an empirical DC clique, in the sense that by design, a nonempty fixed point A meets the first requirement of a DC clique in Definition 1.1 up to a level of statistical significance. The only other possible outcome of the iterative search procedure is a multi-set cycle; this corner case is discussed in Section 2.5.

Remark. Iterative updating using multiple testing was applied by Wilson et al. (2014) in the context of community detection for binary networks. DCM employs a similar search structure in the context of differential correlation. Whereas the method of Wilson et al. relies on a null network distribution that is based on observed node degrees, here we make use of a central limit theorem that does not rely on the structural properties of the observations.

2.4 Residualization

In general, we expect multiple DC cliques to be present in data. The residualization step allows the DCM procedure to search the same dataset many times, avoiding repeat

identification of prior results. This is accomplished by generating new residualized data matrices $\mathbb{X}_1', \mathbb{X}_2'$ after each (non-degenerate) termination of the set update algorithm.

For clarity, let us restrict our attention to $\mathbb{X}_1$ (the same process is applied, separately, to $\mathbb{X}_2$). Suppose the set update procedure converges to a set A, with $|A| = k > 0$. We model the correlation matrix of the variables in A as the combination of a rank-one shared group correlation, characterized by $\Lambda \in \mathbb{R}^k$, and an underlying residual structure $\Omega \in \mathbb{R}^{k \times k}$, so that $\mathbf{R}_{1,A} = \Lambda\Lambda^t + \Omega$. The problem of "factor analysis", or representing correlation matrices by the best lower-rank approximation, is well studied (Harman, 1960), and efficient methods for estimating Λ are readily available.[2] Given an estimate $\hat{\Lambda}$, we generate a new submatrix $\mathbb{X}_{1,A}'$ such that

$$\widehat{\mathrm{cor}}_n\left(\mathbb{X}_{1,A}'\right) = \widehat{\mathbf{R}}_{1,A} - \hat{\Lambda}\hat{\Lambda}^t = \widehat{\Omega}.$$

We then build $\mathbb{X}_1'$ by replacing $\mathbb{X}_{1,A}$ with the residualized data $\mathbb{X}_{1,A}'$ and leaving the rest of the data matrix unaltered. In this way, neither $\mathbb{X}_1'$ nor $\mathbb{X}_2'$ contain the groupwise structure of A captured by Λ. We are then free to apply the DCM procedure to $\mathbb{X}_1'$ and $\mathbb{X}_2'$ and search for secondary empirical DC cliques.

2.5 Special Cases

Minimality: A nonempty fixed point A of the set update procedure has the property that, analogously to Definition 1.1, $H_0(i, A)$ is rejected if and only if $i \in A$. The second condition of Definition 1.1, however, is not guaranteed in general. It is possible that DCM may select a large set that in truth consists of two (or more) disjoint population DC cliques. These cases are well addressed by the residualization step. When a conglomerate estimated DC clique is residualized, the *joint* structure is removed, leaving behind the individual structure of the disjoint DC cliques. Further runs of the DCM algorithm are then able to identify the separate DC cliques.

In extreme cases, the sampled data may be such that the disjoint DC cliques are, by chance, correlated enough to have negligible remaining individual structure after residualization. This renders the multiple DC cliques indistinguishable in the data from a combined DC clique.

Cycles: Under certain conditions, the main search procedure may be caught in a cycle of two or more sets, so that termination to a fixed point never occurs. When the set update procedure oscillates between two sets A_1 and A_2, we restart the search on the intersection $A = A_1 \cap A_2$. Usually, this results in true convergence to a fixed point in the vicinity of the intersection. On occasion, the same oscillation (A_1 to A_2) re-emerges, in which case the overlap $A = A_1 \cap A_2$ is selected as the final output. For this overlap set, $H_0(i, A)$ will

[2] In the DCM software, we use the implementation of Friguet et al. (2012) for the R Statistical Software version and the method of Bishop (2006) for the Matlab version.

be rejected for all $i \in A_1, A_2$. It is not quite a fixed point; nevertheless, it shares many properties of a true DC clique, and we consider it worthy of attention.

Cycles of length greater than two are rarely observed in real or simulated data. However, to protect against longer cycles leading to infinite loops, the algorithm is built with an adjustable maximum iteration limit.

3 Properties of the Test Statistic

We now discuss some properties of the test statistic $\hat{\Delta}(i, A)$ used in the calculation of p-values for the set update procedure.

3.1 Geometric Interpretation

The equation for $\hat{\Delta}(i, A)$ given in (3) expresses the test statistic directly in terms of average differential correlation. However, we may also write $\hat{\Delta}(i, A)$ in an alternate form that yields an informative geometric interpretation. Let $\tilde{\mathbf{U}}_i \in \mathbb{R}^{n_1}$ and $\tilde{\mathbf{V}}_i \in \mathbb{R}^{n_2}$ be the standardized measurements of variable i under sample conditions 1 and 2, respectively; and let

$$\mathbf{W}_1 := \frac{1}{|A|} \sum_{j \in A} \tilde{\mathbf{U}}_j \quad \text{and} \quad \mathbf{W}_2 := \frac{1}{|A|} \sum_{j \in A} \tilde{\mathbf{V}}_j \tag{5}$$

be the vector means of the standardized measurements of the variables in A under each condition. It is easily shown that

$$\frac{1}{|A|} \sum_{j \in A} \widehat{\text{cor}}_n \left(\mathbf{U}_i, \mathbf{U}_j \right) = \mathbf{W}_1^t \tilde{\mathbf{U}}_i = \|\mathbf{W}_1\| \, \widehat{\text{cor}}_n \left(\tilde{\mathbf{U}}_i, \mathbf{W}_1 \right) .$$

Thus,

$$\hat{\Delta}(i, A) = \|\mathbf{W}_1\| \, \widehat{\text{cor}}_n \left(\mathbf{W}_1, \tilde{\mathbf{U}}_i \right) - \|\mathbf{W}_2\| \, \widehat{\text{cor}}_n \left(\mathbf{W}_2, \tilde{\mathbf{V}}_i \right) .$$

Note that the vector $\tilde{\mathbf{U}}_i$ and the vectors $\{\tilde{\mathbf{U}}_j : j \in A\}$ lie on the surface of an $(n_1 - 2)$-dimensional sphere in $\mathbb{R}^{n_1}$, and that $\mathbf{W}_1$ is the geometric center (centroid) of the latter collection. The norm $||\mathbf{W}_1||$ is between 0 and 1; large values of $||\mathbf{W}_1||$ place the centroid closer to the surface of the sphere, indicating that the vectors $\{\tilde{\mathbf{U}}_j : j \in A\}$ are tightly clustered, or equivalently, highly intercorrelated. Thus the quantity $\|\mathbf{W}_1\| \, \widehat{\text{cor}}_n \left(\mathbf{W}_1, \tilde{\mathbf{U}}_i \right)$ weights the similarity of $\mathbf{U}_i$ and the centroid $\mathbf{W}_1$ according to the overall similarity of the vectors $\{\tilde{\mathbf{U}}_j : j \in A\}$. Similar remarks apply to $\{\tilde{\mathbf{V}}_j : j \in A\}$ and $\mathbf{W}_2$. The statistic $\hat{\Delta}(i, A)$ is the difference of the summary measures in conditions 1 and 2.

Figure 3 gives a simple two-dimensional representation of the geometric picture discussed above. In condition 1, $\mathbf{U}_i$ is not strongly correlated with $\mathbf{W}_1$, but $\|\mathbf{W}_1\|$ is large because the vectors indexed by A are tightly clustered. In condition 2, $\mathbf{V}_i$ is strongly

correlated with $\mathbf{W}_2$, but $\|\mathbf{W}_2\|$ is small because the vectors indexed by A are not tightly clustered. In this example, $\hat{\Delta}(i, A)$ is close to zero, and we would likely conclude no differential correlation is present.

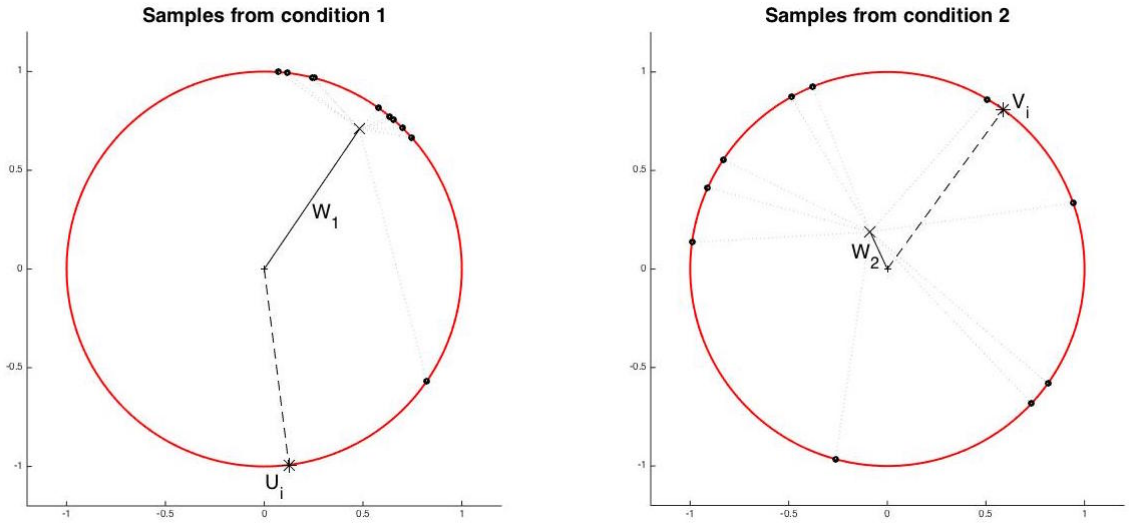

Figure 3: Geometric representation of data in two dimensions. Unlabelled points represent the standardized data in group A.

3.2 Asymptotic distribution of $\hat{\Delta}(i, A)$

We now discuss the asymptotic distribution of $\hat{\Delta}(i, A)$, from which the p-values used in Section 2.3 are derived. First, we make note of a classical result concerning sample correlations.

Theorem 1. (Steiger and Hakstian, 1982) *Let* $\mathbf{R}$ *be a* $p \times p$ *correlation matrix, and* $\widehat{\mathbf{R}}$ *the corresponding sample correlation matrix based on* n *i.i.d. samples of* p*-variate data with finite 4th moment. Let* $\mathbf{P}$ *and* $\widehat{\mathbf{P}}$ *be the vectorized versions of the matrices, of dimension* $p^2 \times 1$*. Then, as* n *tends to infinity*

$$\sqrt{n}\left(\widehat{\mathbf{P}} - \mathbf{P}\right) \Rightarrow \mathcal{N}_{p^2}\left(\mathbf{0}, \mathbf{\Sigma}\right)$$

where $\mathbf{\Sigma}$ *is a* $p^2 \times p^2$ *covariance matrix for which a general form is given equations (3.1-3.5) in Browne and Shapiro (1986).*

Using Theorem 1 one may evaluate the asymptotic distribution of $\hat{\Delta}(i, A)$, which is a function of $\mathbf{P}$ and $\widehat{\mathbf{P}}$.

Corollary 1.1. *Let A be a fixed index set and let $\hat{\Delta}(i,A)$ be defined as in (3), with sample correlation matrices $\widehat{\mathbf{R}}_1$ and $\widehat{\mathbf{R}}_2$ based on n_1 and n_2 independent samples from distributions F_1 and F_2 respectively. Let $\sigma_0^2(i,A) := \mathrm{Var}\left(\hat{\Delta}(i,A) \,|\, H_0\right)$, where H_0 is the null hypothesis in (2). Then, under H_0,*

$$\frac{\hat{\Delta}(i,A)}{\sigma_0(i,A)} \Rightarrow \mathcal{N}(0,1) \tag{6}$$

as $\min(n_1,n_2) \to \infty$.

Proof: For clarity, we first examine only one "half" of $\hat{\Delta}(i,A)$. Let

$$\bar{r}_1(i,A) \,=\, \frac{1}{|A|}\sum_{j\in A}\left(\widehat{\mathbf{R}}_1\right)_{ij} \quad \text{and} \quad \bar{\rho}_1(i,A) \,=\, \frac{1}{|A|}\sum_{j\in A}(\mathbf{R}_1)_{ij}\,. \tag{7}$$

Note that $\bar{r}_1(i,A)$ is a linear function of $\widehat{\mathbf{R}}_1$ and that $\bar{\rho}_1(i,A)$ is the same function applied to the population correlation matrix $\mathbf{R}_1$. It follows from Theorem 1 that

$$\sqrt{n_1}\left(\frac{\bar{r}_1(i,A)-\bar{\rho}_1(i,A)}{\tau_1^2(i,A)}\right) \Rightarrow \mathcal{N}(0,1)\,.$$

with $\tau_1^2(i,A) := \mathrm{Var}\left(\sqrt{n_1}\,\bar{r}_1(i,A)\right)$, which has a finite limiting value that can be expressed as the mean of appropriate elements of the covariance matrix $\mathbf{\Sigma}$ in the theorem. To apply this result for the full test statistic, we note that $\hat{\Delta}(i,A) = \bar{r}_1(i,A) - \bar{r}_2(i,A)$. Samples from F_1 are independent of those from F_2, so $\bar{r}_1(i,A)$ is independent of $\bar{r}_2(i,A)$, and thus $\hat{\Delta}(i,A)$ is asymptotically normal.

Under the null hypothesis in (2), $\bar{\rho}_1(i,A) = \bar{\rho}_2(i,A)$, and therefore the mean of the limiting distribution of $\hat{\Delta}(i,A)$ is 0. The variance of $\hat{\Delta}(i,A)$ can be expressed as the weighted sum

$$\sigma_0^2(i,A) \,=\, \frac{\tau_1^2(i,A)}{n_1} + \frac{\tau_2^2(i,A)}{n_2}\,.$$

□

In practice, the variances $\tau_1^2(i,A)$ and $\tau_2^2(i,A)$ are not known. Thus, we make use of the results in Steiger and Hakstian (1982) to identify the asymptotic variance of $\bar{r}_1(i,A)$, and a consistent estimator $\hat{\tau}_1$.

Lemma 1.1. *Let r_{ij} be the sample correlation of $\mathbf{U}_i$ and $\mathbf{U}_j$, and let $r_A := \frac{1}{|A|}\sum_{j\in A} r_{ij}$. For $\ell = 1,2,...,n_1$ let*

$$W_\ell \,:=\, \frac{1}{|A|}\sum_{j\in A}\tilde{\mathbf{U}}_{j\ell}\,, \quad \textit{and} \quad Y_\ell \,:=\, \frac{1}{|A|}\sum_{j\in A} r_{ij}\tilde{U}_{j\ell}^2\,.$$

and define

$$\hat{\tau}_1 \;=\; \frac{1}{n_1}\sum_{\ell=1}^{n_1}\left\{\frac{r_A^2}{4}\,\tilde{U}_{i\ell}^4 \;-\; r_A W_\ell\,\tilde{U}_{i\ell}^3 \;+\; \left(\frac{r_A Y_\ell}{2} \;+\; W_\ell^2\right)\tilde{U}_{i\ell}^2 \;-\; W_\ell Y_\ell\,\tilde{U}_{i\ell} \;+\; \frac{Y_\ell^2}{4}\right\}.$$

If $\bar{r}_1(i,A)$ is defined as in (7)*, then under any sampling distribution satisfying the assumptions of Theorem 1, we have*

$$\frac{\hat{\tau}_1}{\tau_1} \;\xrightarrow{p}\; 1 \quad \textit{as} \quad n_1 \to \infty\,,$$

where $\tau_1 = \mathrm{Var}\left(\sqrt{n_1}\,\bar{r}_1(i,A)\right)$.

Lemma 1.1 is proven in Appendix A. The results of this lemma allow us to make an efficient and straightforward calculation to estimate the variances of $\bar{r}_1(i,A)$ and $\bar{r}_2(i,A)$ for every $i \in [p]$, used in the testing step of the DCM algorithm. Note that regardless of the size of A, the derived formula for $\hat{\tau}_1$ involves basic operations on only three $p \times n_1$ vectors: $\mathbf{W}, \mathbf{Y}$ and $\tilde{\mathbf{U}}_i$. Such simplification is important, since in practice the estimates $\hat{\tau}_1$ and $\hat{\tau}_2$ must be calculated for *every* variable $i \in [p]$ and for multiple iterative steps of A.

Remark. The results of Corollary 1.1 and Lemma 1.1 apply to variable sets of fixed cardinality ($|A| = k$) as n_1 and n_2 tend to infinity. In practice, one may encounter variables sets for which $k > n_1, n_2$. Simulations suggest that the DCM algorithm still identifies DC cliques with high success and controls false discovery in such settings even when the cardinality of $|A|$ greatly exceeds the sample size.

4 Simulation Study

To test the DCM method against possible alternatives, we implemented a simple study of performance on simulated data. We created artificial datasets containing a single DC clique and compared the results of several methods to the known truth. Although the simulated setting is not a perfect representation of real data situations, it readily illustrates the advantages of DCM as opposed to adaptations of existing methods.

4.1 Methods implemented

In the absence of comparable methods, we adapted several existing methods to search for DC cliques. These adaptations are described below.

Mining a single correlation matrix (WGCNA). We implemented the Weighted Gene Co-Expression Network Analysis (WGCNA) method of Langfelder and Horvath (2008) for clustering correlation matrices of gene expression data. This method is a hybrid analysis, which mines for clusters (or "modules") in a single correlation matrix, then tests each

module for differential *expression*. Thus, although the WGCNA method involves both differential and second-order elements, it is not designed to search for DC cliques or similar structures. For the purposes of this simulation study, we applied WGCNA to samples from class 1 only. We then tested the output module for differential correlation via a standard t-test over sample correlations in classes 1 and 2. In this way, we attempted to only select variable sets exhibiting differential correlation, even though WGCNA does not naturally identify modules with this property.

Estimation of differential correlation matrices (D-EST and FISH). We implemented the method of Cai and Zhang (2014) to estimate $\mathbf{D} = \mathbf{R}_1 - \mathbf{R}_2$ by the suggested estimator $\widehat{\mathbf{D}}$. To search for the seeded DC clique, we converted $\widehat{\mathbf{D}}$ to a $[0, 1]$ scale (i.e. a dissimilarity matrix $\mathbf{D}^*$), then performed hierarchical clustering on $\mathbf{D}^*$. We extracted the first cluster in the dendrogram, starting at the bottom, of size less than or equal to the size of the seeded DC clique. This approach necessarily returns a nonempty variable set of approximately correct size; in non-simulated settings, one would not know the size of true DC cliques.

In a similar fashion, we calculated

$$\widehat{\mathbf{D}}_{Fisher} = (n_1 - 3)^{1/2}\, \varphi\left(\widehat{\mathbf{R}}_1\right) - (n_2 - 3)^{1/2}\, \varphi\left(\widehat{\mathbf{R}}_2\right)$$

where φ is the (element-wise) Fisher transformation of sample correlations. Note that $\widehat{\mathbf{D}}_{Fisher}$ is not, in fact, an estimator of $\widehat{\mathbf{R}}_1 - \widehat{\mathbf{R}}_2$, but it is analogous to $\widehat{\mathbf{D}}$ in that large entries in general correspond to large correlation differences. As with $\widehat{\mathbf{D}}$, we converted $\widehat{\mathbf{D}}_{Fisher}$ to a dissimilarity matrix and extracted a cluster from a hierarchical dendrogram.

Detection of isolated changes in correlation structure (DCP). We applied the Differential Correlation Profile (DCP) method of Liu et al. (2010). This method takes a permutation approach to variable selection. Each variable $i \in [p]$ is assigned a measure Ψ_i of overall differential correlation. The samples are then randomly permuted many times, each time calculating new measures Ψ_i', and significant variables are chosen by comparing Ψ_i to the permutated values. This approach identifies a list of individual differentially correlated variables, rather than a united set; for the purposes of this study, we considered the collection of selected variables to be the estimated DC clique.

Alternate testing step (CLASSIC). Finally, we implemented an altered version of DCM that uses the same algorithmic framework but makes use of classical results to test for differential correlation. Instead of testing $H_0(i, A) : \hat{\Delta}(i, A) = 0$ in the DCM algorithm (as in Section 2.3), we test $H_0(i, A) : R_1(i, A) = R_2(i, A)$, or equality of the vectors of correlations between i and elements in A. We adapted the testing method of Larntz and Perlman (1985) for a one-sided alternate hypothesis; namely, $H_a(i, A) : \max\{R_1(i, A)\} > \max\{R_2(i, A)\}$.

4.2 Simulated Data

We simulated data with a single embedded DC clique, consistent with Definition 1.1. Our study varied the following parameters: size of the DC clique (k), total number of variables

(p), strength of the correlations (ρ_1 and ρ_2), samples sizes of the two datasets (n_1 and n_2), and background data. We considered three background types: uncorrelated, positive, and real data. To elaborate, let $\mathbf{S}_1$ and $\mathbf{S}_2$ be matrices that are 1 on the diagonal, ρ_1 or ρ_2 (respectively) on the off-diagonal for indices 1 to k, and 0 otherwise. Then, the three types of data simulations were:

Uncorrelated Gaussian. The data was generated from multivariate Normal distributions with correlations $\mathbf{R}_1 = \mathbf{S}_1, \mathbf{R}_2 = \mathbf{S}_2$.

Positively correlated Gaussian. The correlation matrices of the data were $\mathbf{R}_1 = \mathbf{S}_1 + \rho_1/3$, $\mathbf{R}_2 = \mathbf{S}_2 + \rho_1/3$. That is, the correlation matrix was boosted everywhere (except on the diagonals) in equal part for both datasets.

Real data. The dataset was derived from a real-world data source[3] with samples randomly assigned one of two sample classes. By adding a common vector to the first k rows of the data matrix, we induced further correlation ρ_1 and ρ_2 in part the data.

We found that all methods behaved similarly with regard to changes in sample sizes n_1, n_2 and clique size k (relative to p). Here, we present only the results regarding ρ_1 vs. ρ_2 and different background types, to illustrate key differences in performance between methods. By default, the other parameters were set to be $n_1 = n_2 = 100$, $k = 100$, and $p = 1000$.

4.3 Results

We applied the five proposed methods (DCM, WGCNA, D-EST, FISH, DCP, and CLASSIC) to a variety of values of ρ_1 and ρ_2 for each of the three background data types. For each (ρ_1, ρ_2), we ran 100 random trials and averaged the results. The success of the methods was measured by the false positive rate (FPR), the percentage of variables in a selected set that were not in the seeded DC clique, and the true discovery rate (TDR), the percentage of detected variables from the true DC clique. That is, if B was the output variable set of a procedure and $A = (1, \ldots, k)$ was the embedded DC clique, then

$$\text{False Positive Rate} = \frac{|B \setminus A|}{|B|} \quad \text{and} \quad \text{False Negative Rate} = \frac{|A \setminus B|}{|A|}.$$

Figure 4 shows the False Positive rate for each method as the seeded differential correlation gets stronger ($\rho_1 - \rho_2$ grows), with other parameters fixed at default values, for each background type. Figure 5 shows the same for False Negative rate. In both cases,

[3]The Cancer Genome Atlas; Luminal A tumor subtype

values close to zero are desirable, as they represent low error in identifying the true seeded DC clique.

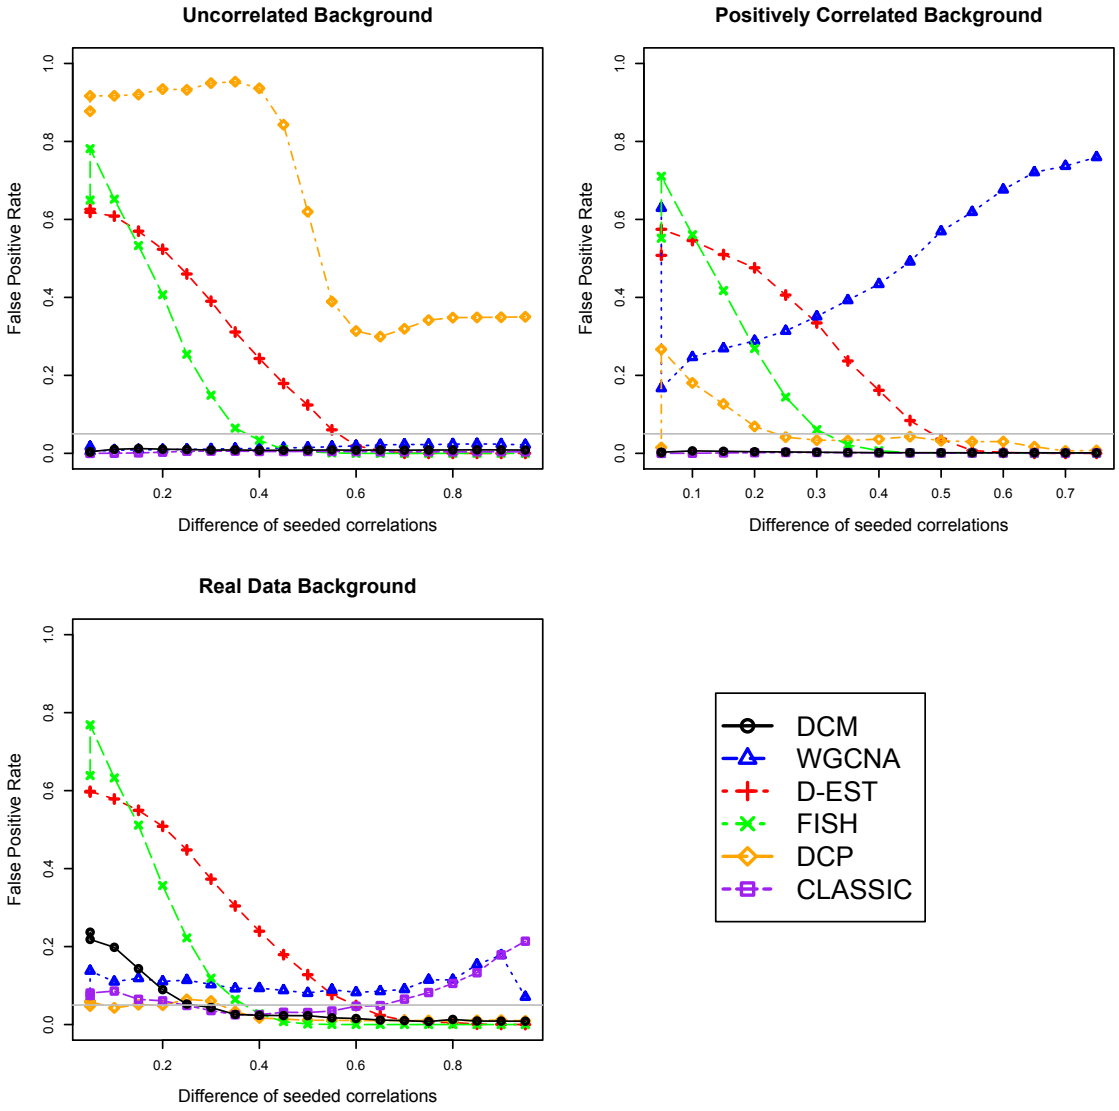

Figure 4: Average False Positive rate for each data type. ($n1 = n2 = 100, k = 100, p = 1000$).

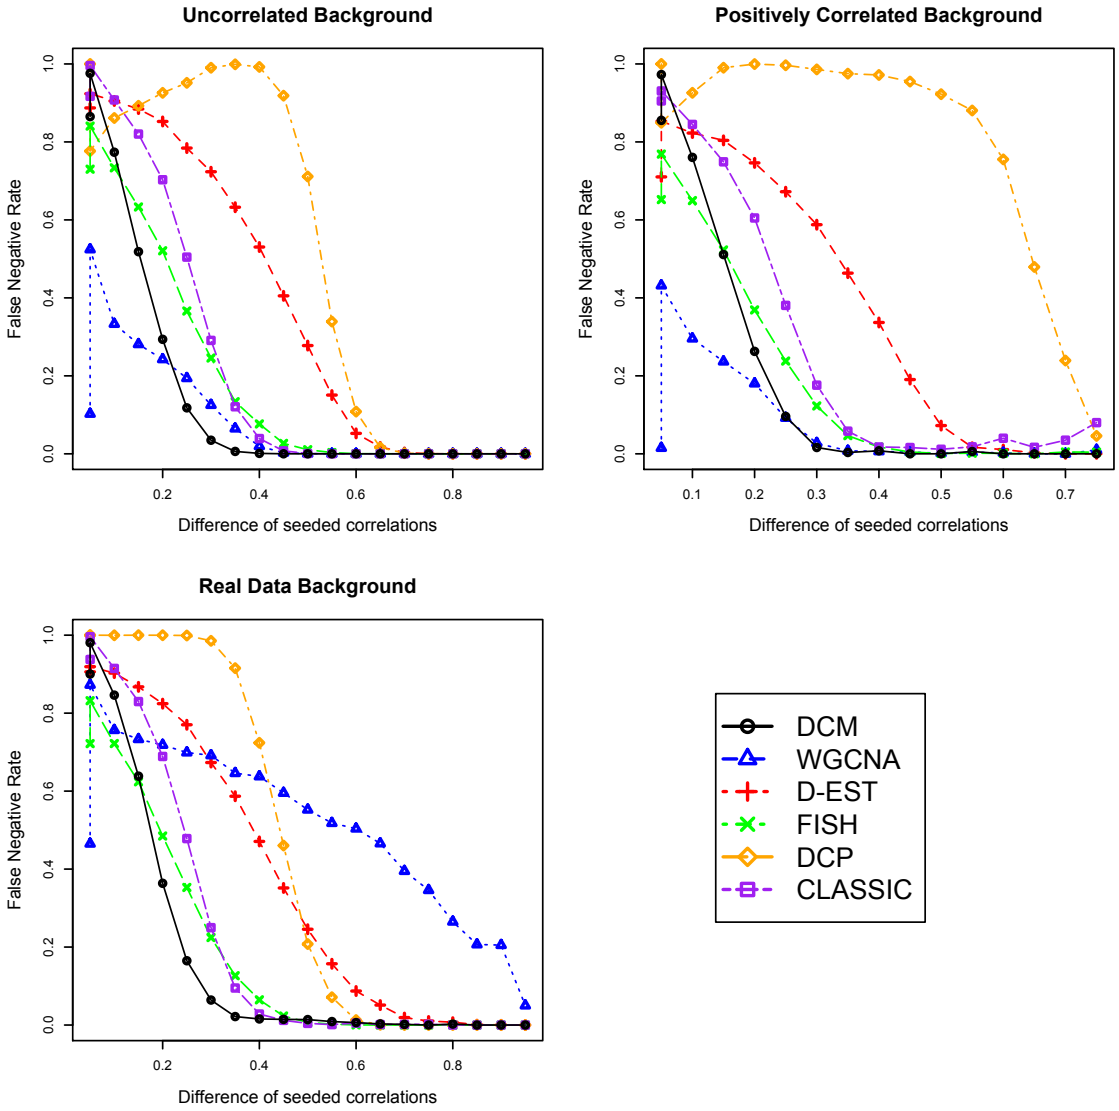

Figure 5: Average False Negative rate for each data type. ($n1 = n2 = 100, k = 100, p = 1000$).

We note that DCM controls the rate of false positives in all cases, except for some error when there is very low signal size in the real data case, which may be due to actual signal being present in the randomized real data. DCM also begins to reliably detect DC cliques at a lower signal size - around a correlation difference of 0.2 - in every setting except for the simple uncorrelated background setting, where WGCNA has slightly better performance.

When the background is highly correlated, but not differentially correlated, WGCNA is not able to control the false positive rate. This is because WGCNA is designed to be a single-dataset mining method. When the background data shows strong correlation, WGCNA correctly identifies a large correlated module in the first dataset. Even if these modules could be tested for differential correlation, there is no method inherent to WGCNA for extracting the DC clique from surrounding correlated (but not differentially correlated) data.

The D-EST and FISH methods behave as expected; because our approach necessarily returns a nonempty variable set, the false positive rate is high for small signal. However, even if the false positives were perfectly controlled in some way, these methods show a higher false negative rate than DCM. Similarly, the adjusted DCM algorithm with the classical test (CLASSIC) controls error, but is less powerful than our method.

Finally, the DCP method vastly overselects variables in the uncorrelated background case. This is likely because the mutual behavior of the variables in A induces some correlation structure in the background variables; Figure 1 illustrates this phenomenon, as there is some pattern in the cross correlation between variables in B and A. This result emphasizes the danger of the common approach of looking for isolated changes in correlation structure of individual variables, rather than searching for DC cliques: vestigial correlation patterns may be misleading. (When the background has structure, as in positively correlated or real data, the induced pattern is overshadowed, so the high false positive rate of DCP does not occur.)

We wish to note in particular the behavior of these methods when the correlation difference is zero. These simulations represent settings where correlation is present in both datasets, but the structure is identical across datasets rather than differential. We expect such scenarios to arise commonly in real data, where groups of variables may be universally correlated without regard to the particular sample conditions being studied. It is important that a search procedure does not identify nonexistent DC cliques in these situations. Figure 6 shows the average size of selected sets as a function of the strength of the *non-differential* correlation. Values at zero represent cases where the method successfully avoided identifying the misleading correlation as a DC clique; large values represent the discovery of a false positive set.

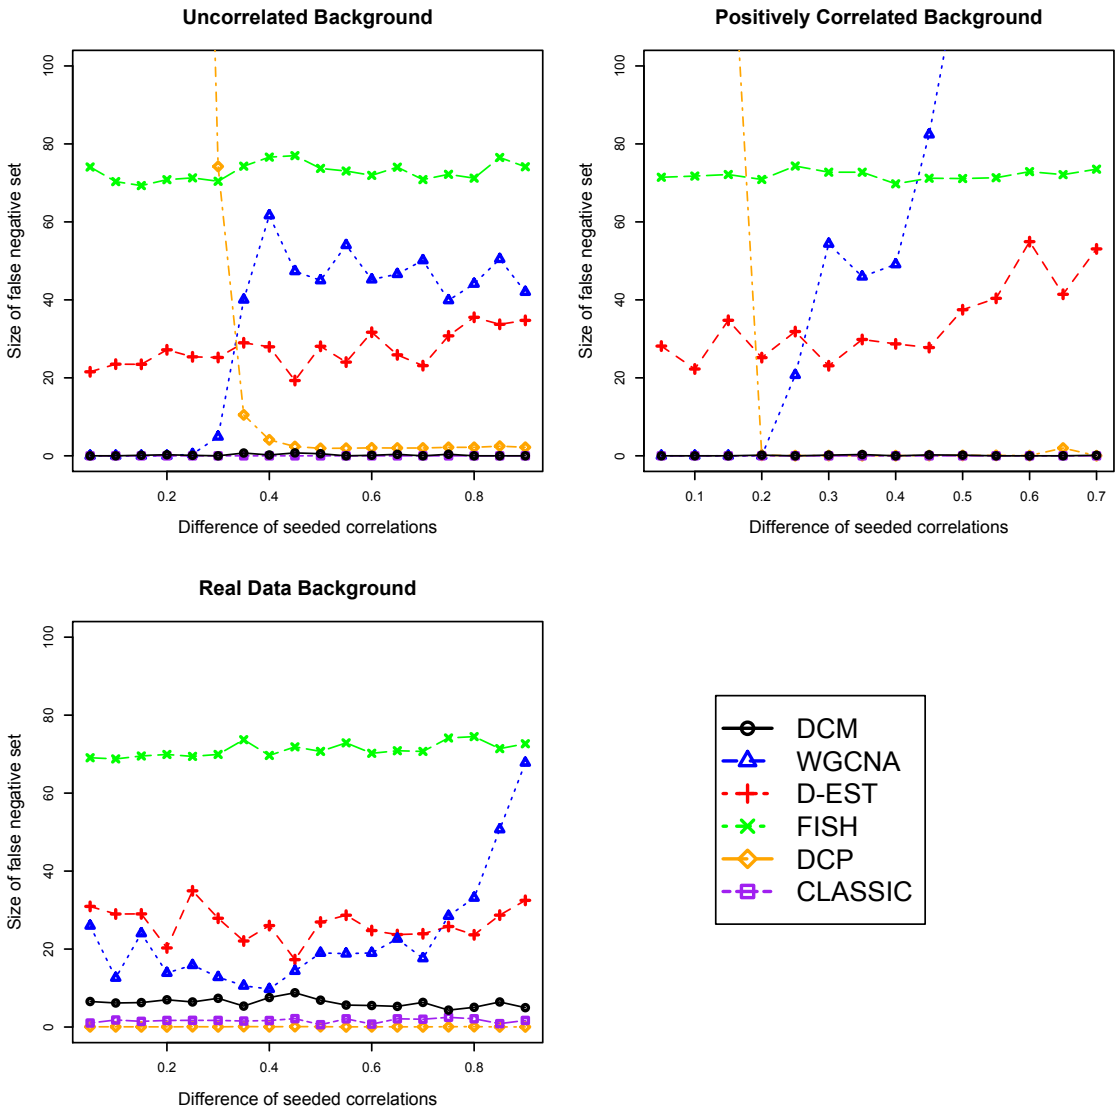

Figure 6: Average size of set when no DC clique is present, for each data type. ($n1 = n2 = 100, k = 100, p = 1000$).

It is clear that all methods except DCM and CLASSIC are prone to error when non-differential correlation is present. This striking difference in performance illustrates the ways in which existing methods are simply not designed for the specific case of DC cliques. Only DCM-type algorithms include a testing-based element, which allow them to dismiss observed correlation for which there is not enough evidence of differential behavior. (In the

real data case, DCM sometimes selects small nonempty sets in spite of the lack of seeded DC clique; again, this is likely due to the fact that permuting the samples in real data will sometimes produce signal.)

4.4 A note about computation

Of the methods tested, only DCM, CLASSIC, and WGCNA are computationally practical for large p ($\sim 10^5$ or more). The methods FISH and D-EST have memory demands on the order of p^2, as they are based upon the estimation of the full p by p dissimilarity matrices, $\widehat{\mathbf{D}}_{Fisher}$ and $\widehat{\mathbf{D}}$. Permutation-based methods such as DCP are even more infeasible, since they require the computation of a p by p correlation matrix for each of many permutations.

5 Data Analyses

5.1 TCGA

We now expand upon the real data application referenced in Figure 1. Recall that we applied the DCM procedure to samples from two pre-determined breast cancer subtypes: Her-2 and Luminal B. A total of 18 empirical DC cliques (more correlated in Her-2 than in Luminal B) were discovered, ranging in size from 13 to 165 genes. To illustrate how this information may be useful to genomic research, we briefly discuss one of the discovered gene sets. The set of interest contained 46 genes, listed alphabetically in Table 1. These genes are found to be highly associated with immune response, particularly the HLA (Human Leukocyte Antigen) gene class, represented by six of the genes in the set. Researchers are interested in understanding how and why some cancer subtypes trigger immune response while others do not. For example, Iglesia et al. (2014) showed that prognosis was improved for patients with Her-2 and basal-like subtypes showing higher immunoreactive response. Further exploration of DC cliques such as the one in Table 1 may further understanding of the gene interactions that drive immune response.

Table 1: Genes selected in empirical DC Clique for Her2 vs. Luminal B samples.

AGER	amt	APOL1	ARPC4	B2M	BATF2	BTN3A2
BTN3A3	C19orf38	calml4	CCDC146	CHKB-CPT1B	echdc1	ETV7
EXOSC10	FBXO6	GBP1	GBP4	GJD3	gnb3	**HLA-A**
HLA-B	**HLA-C**	**HLA-E**	**HLA-F**	**HLA-H**	HSH2D	IDO1
IL15	Irf1	LOC115110	LOC400759	LOC91316	micB	Myo15b
OASL	PILRB	Rec8	Rufy4	SAMD9L	SEC31B	STAT1
tap1	Tapbp	TTLL3	TXNDC6	Ube2l6	Zbp1	

5.2 The Human Connectome Project

The Human Connectome Project is a multi-institutional venture aimed at mapping functional connections between parts of the human brain. The project has collected vast amounts of brain scan data, all of which is publicly available to researchers online at www.humanconnectome.org. In this analysis, we made use of a dataset from the "500 Subjects MR" data release, which consists of functional magnetic resonance imaging (fMRI) brain scans for 542 healthy adult subjects. Participants performed a variety of tasks during the MR scan, designed to isolate certain types of brain functionality. Activation levels were recorded over time for ∼30,000 voxels (3D coordinate locations in the brain's white matter interior) and ∼60,000 greyordinates (indexed locations over the grey matter brain surface).

In this paper, we applied DCM to data from a single subject.[4] We compared two task categories:

> *Language-based tasks:* During the scan, subjects were told brief stories and asked to answer questions after each one about what they were told.
>
> *Motor-based tasks:* Subjects were attached to motion sensors at the hands, feet, and tongue. They were then asked to move one appendage at a time, in blocks of repetitions.

DCM to searched amongst 91,282 brain locations (or nodes) for DC cliques that exhibit more correlation over time during language tasks than during motor tasks.

The first empirical DC clique selected by DCM contained 913 nodes located on the cortical surface. These nodes, or "greyordinates", are visualized as points on the smoothed exterior of the brain in Figure 7. The clear locational pattern in the nodes - despite the fact that the analysis did not take location into account - is striking. Additionally, the empirical DC clique in Figure 7 includes a concentrated group in the rear of the left cortex. This general brain region is known to be specifically associated with language processing and auditory input (Wernicke's Area, see Wang et al. (2015)).

We also studied two other artifacts of the data for comparison. First, we identified the 1000 nodes exhibiting the strongest differential first-order behavior. These show higher mean activation during the language tasks than during the motor tasks, as measured by standard two-sample t-tests. We saw a clear grouping of nodes in the right frontal lobe. This pattern is unsurprising and appears in many studies of brain functionality that examine differential activation for language processing (Voets et al., 2006). This basic first-order analysis suggests that differential correlation is not redundant. None of the empirical DC cliques selected by DCM show high frontal lobe concentration; instead, they exhibit "trail-like" patterns such as the ones shown in Figure 7.

Second, we identified 1000 nodes found to be highly correlated over time for the language task data, irrespective of their behavior in the motor task data. These nodes were

[4]Subject #101006, a 35-year-old female.

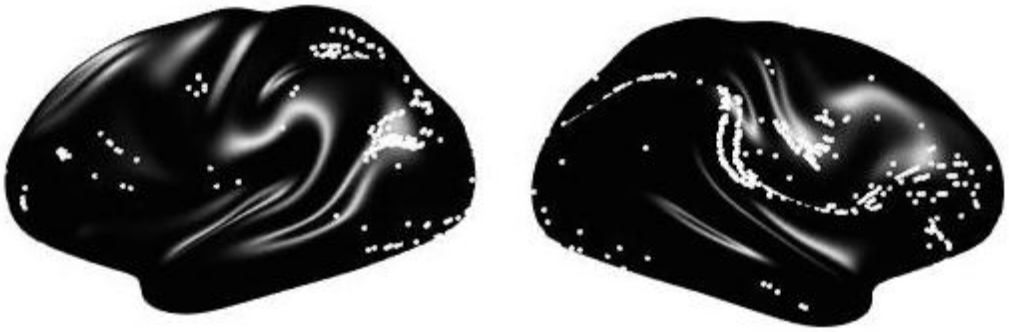

Figure 7: Brain locations of DC clique for languages tasks versus motor tasks.

observed to be very tightly grouped in the interior left hemisphere. This is likely due to the nature of data measurement: fMRI brain scans measure oxygen flow in the brain, so measurements for adjacent regions tend to "blur" and show high artificial correlation (Derado et al., 2010). In this case, the same node set is also highly correlated during motor tasks, suggesting that it is likely a byproduct of data collection. Even if this node set does represent a meaningful result - regions, perhaps, that are universally correlated regardless of task - it is not differential.

This example illustrates the advantage of taking a differential approach like DCM. Effects due to fMRI-driven spatial correlation or strong universal correlation can drown out signal that is truly specific to a particular sample condition. By comparing language tasks to the similar but distinct condition of motor tasks, we are able to isolate signals that are unique to language processing. The fact that the identified DC cliques show emergent locational patterns suggests that DCM is capturing a true facet of the data rather than arbitrary correlation. Since this output is unique in form, while maintaining some consistency with known brain functionality, we believe it merits further scientific investigation.

6 Conclusion

In this paper, we have introduced a new statistical method, DCM, to identify differentially correlated variable sets from observed data. The DCM algorithm has been shown to be built on statistical principles with a theoretical basis, and to perform accurately and efficiently in a variety of settings. Unlike existing methods, DCM is specifically built to discover DC cliques, and its underlying testing process controls error. Additionally, the DCM software can be run on extremely high dimensional data without large memory demands or long runtimes. Preliminary data analysis results in the application areas of both gene expression data and brain activation are encouraging.

Code for public use of DCM is freely available at http://kbodwin.web.unc.edu/software/.

Acknowledgements

The authors wish to thank Yin Xia, Andrey Shabalin, Katherine Hoadley, and Kimberly D.T. Stachenfeld for their contributions.

A Variance Estimator

A.1 Proof of Lemma 1.1

We wish to show that if $\bar{r}_1(i, A)$ is defined as in (7), for fixed $i \in [p]$ and $A \subset [p]$, then

$$\hat{\tau}_1 \;=\; \frac{1}{n_1}\sum_{\ell=1}^{n_1}\left\{\frac{r_A^2}{4}\,\tilde{U}_{i\ell}^4 \;-\; r_A W_\ell\,\tilde{U}_{i\ell}^3 \;+\; \left(\frac{r_A Y_\ell}{2} \;+\; W_\ell^2\right)\tilde{U}_{i\ell}^2 \;-\; W_\ell Y_\ell\,\tilde{U}_{i\ell} \;+\; \frac{Y_\ell^2}{4}\right\}.$$

is equal to the consistent estimator for the variance of $\bar{r}_1(i, A)$ given in Steiger and Hakstian (1982). Define

$$r_{ijkh} \;:=\; \sum_{\ell=1}^{n_1}\tilde{U}_{i\ell}\tilde{U}_{j\ell}\tilde{U}_{k\ell}\tilde{U}_{h\ell}\,.$$

Then, for the variance of an average of sample correlations sample correlations, $\tau_1 = \frac{1}{|A|^2}\sum_{j,k\in A}\mathrm{Cov}\,(r_{ij}, r_{ik})$, equation (5.1) of Steiger and Hakstian (1982) gives the consistent estimator

$$\begin{aligned}
\hat{\tau}_1 \;&=\; \frac{1}{|A|^2}\sum_{j,k\in A}\Bigg[r_{iijk} \;+\; \frac{1}{4}r_{ij}r_{ik}(r_{iiii} + r_{iijj} + r_{iikk} + r_{jjkk}) \\
&\quad -\; \frac{1}{2}r_{ij}(r_{iiik} + r_{ijjk}) \;-\; \frac{1}{2}r_{ik}(r_{iiij} + r_{ijkk})\Bigg] \\
&=\; \frac{1}{|A|^2}\sum_{j,k\in A} r_{iijk} \;+\; \frac{1}{4}r_A^2 r_{iiii} \;+\; \frac{1}{2|A|}\sum_{j\in A} r_{ij}r_A r_{iijj} \;+\; \frac{1}{4|A|^2}\sum_{j,k\in A} r_{ij}r_{ik}r_{jjkk} \\
&\quad -\; \frac{1}{|A|}\sum_{j\in A} r_A r_{iiij} \;-\; \frac{1}{|A|^2}\sum_{j,k\in A} r_{ij}r_{ijjk}
\end{aligned}$$

Using the notation of Section 3.2, we can write

$$\frac{1}{|A|^2}\sum_{j,k\in A} r_{iijk} \;=\; \frac{1}{n_1}\sum_{\ell=1}^{n_1}\left(\frac{1}{|A|^2}\sum_{j,k\in A}\tilde{U}_{j\ell}\tilde{U}_{k\ell}\right)\tilde{U}_{i\ell}^2 \;=\; \frac{1}{n_1}\left(\mathbf{W}^{\circ 2}\right)^t\,\tilde{\mathbf{U}}_i^{\circ 2}\,.$$

$$\frac{1}{|A|}\sum_{j\in A} r_{ij}r_{iijj} \;=\; \frac{1}{n_1}\sum_{\ell=1}^{n_1}\left(\frac{1}{|A|}\sum_{j\in A} r_{ij}\tilde{U}_{j\ell}^2\right)\tilde{U}_{i\ell}^2 \;=\; \frac{1}{n_1}\mathbf{Y}^t\,\tilde{\mathbf{U}}_i^{\circ 2}\,.$$

$$\frac{1}{|A|^2}\sum_{j,k\in A} r_{ij}r_{ik}r_{jjkk} \;=\; \frac{1}{n_1}\sum_{\ell=1}^{n_1}\left(\frac{1}{|A|}\sum_{j\in A} r_{ij}\tilde{U}_{j\ell}^{\circ 2}\right)^2 \;=\; \frac{1}{n_1}\mathbf{Y}^t\,\mathbf{Y}\,.$$

$$\frac{1}{|A|^2}\sum_{j,k\in A} r_{ij}r_{ijjk} \;=\; \frac{1}{n_1}\sum_{\ell=1}^{n_1}\left(\frac{1}{|A|}\sum_{j\in A} r_{ij}\tilde{U}_{j\ell}^2\right)\left(\frac{1}{|A|}\sum_{k\in A}\tilde{U}_{k\ell}\right)\tilde{U}_{i\ell} \;=\; \frac{1}{n_1}\left(\mathbf{W}\cdot\mathbf{Y}\right)^t\,\tilde{\mathbf{U}}_i\,.$$

Thus, we can write $\hat{\tau}_1$ as a simple inner products of the vectors $\mathbf{Y}, \mathbf{W}$, and $\tilde{\mathbf{U}}_i$,

$$n_1\hat{\tau}_1 \;=\; \frac{1}{4}r_A^2\mathbf{1}^t\tilde{\mathbf{U}}_i^{\circ 4} \;+\; r_A\left[\frac{1}{2}\mathbf{Y}^t\,\tilde{\mathbf{U}}_i^{\circ 2} \;-\; \mathbf{W}^t\tilde{\mathbf{U}}_i^{\circ 3}\right] \;+\; \left(\mathbf{W}^{\circ 2}\right)^t\,\tilde{\mathbf{U}}_i^{\circ 2} \;-\; \left(\mathbf{W}\cdot\mathbf{Y}\right)^t\,\tilde{\mathbf{U}}_i \;+\; \frac{1}{4}\mathbf{1}^t\,\mathbf{Y}^{\circ 2}$$

$$\hat{\tau}_1 \;=\; \frac{1}{n_1}\sum_{\ell=1}^{n_1}\left\{\frac{r_A^2}{4}\,\tilde{U}_{i\ell}^4 \;-\; r_A W_\ell\,\tilde{U}_{i\ell}^3 \;+\; \left(\frac{r_A Y_\ell}{2} \;+\; W_\ell^2\right)\tilde{U}_{i\ell}^2 \;-\; W_\ell Y_\ell\,\tilde{U}_{i\ell} \;+\; \frac{Y_\ell^2}{4}\right\}\,.$$

□ *Lemma 1.1*

Remark. We note that although the estimator $\hat{\tau}_1$ is consistent for a very general set of sampling distributions, it may in some cases converge slowly. For very small sample sizes, we find the estimator to be negatively biased; that is, tests involving this estimator may be anticonservative. Although the full DCM procedure appears in simulations to control false positive rate even for small sample sizes, we caution against its use when $\min(n_1, n_2) < 30$.

B Detailed Pseudocode

Algorithm 1 Initial Search Procedure

procedure INITDCM($\mathbf{X}_1, \mathbf{X}_2, k$) ▷ Target output size k.
  $\mathbf{F}_1, \mathbf{F}_2 \leftarrow$ Fisher transformed correlation matrices of $\mathbf{X}_1, \mathbf{X}_2$.
  $B \leftarrow$ index set of size k chosen uniformly at random
  repeat
    $S = \sum_{i,j \in B} (\mathbf{F}_1 - \mathbf{F}_2)_{ij}$
    for a in B, r in B^C **do** ▷ Possible swaps
      $S_{ar} = \sum_{i,j \in B \cup \{r\} \setminus \{a\}} (\mathbf{F}_1 - \mathbf{F}_2)_{ij}$
    end for
    $a^*, r^* \leftarrow$ maximizers of S_{ar} subject to $S_{a^* r^*} > S$.
    $B \leftarrow B \cup \{r^*\} \setminus \{a^*\}$ ▷ Best swap
  until no such a^*, r^* exist
return B
end procedure

Algorithm 2 Core Search Algorithm

procedure DCM($\mathbf{X}_1, \mathbf{X}_2, A$) ▷ Initial index set A
  $A_{prev} \leftarrow \emptyset$
  $cycle \leftarrow 0$
  repeat
    $t = 1$ or 2 ▷ Sample class label
    $m_t \leftarrow \text{mean}\left(\{\mathbf{X}_t\}_A\right)$
    for i in $1, \cdots, p$ **do**
      $r_{ti} \leftarrow \widehat{\text{cor}}\left(X_{ti}, m_t\right)$
      $\hat{T}_i \leftarrow m_1 r_{1i} - m_2 r_{2i}$ ▷ Sample test statistic.
      $H_{0,i} : T_i = 0$
      $p_i \leftarrow P\left(T_i \geq \hat{T}_i \,|\, H_{0,i}\right)$ ▷ DC variable p-values
    end for
    $A_{next} = \{i \,:\, H_{0,i}$ rejected by FDR controlled multiple testing$\}$
    if $A_{prev} = A_{next}$ **then** ▷ Check for cycles.
      $A_{prev} \leftarrow A$
      $A \leftarrow A \cap A_{next}$
      $cycle+ = 1$
    else ▷ Update sets.
      $A_{prev} \leftarrow A$
      $A \leftarrow A_{next}$
    end if
  until $A_{next} = A_{prev}$ or $A_{next} = \emptyset$ or $cycle = 2$
return A
end procedure

References

Anderson, T. (1959). *An Introduction to Multivariate Statistical Analysis.* Wiley-Interscience.

Bassi, F. and Hero, A. (2012). Large scale correlation detection. *Proc. of the IEEE International Symposium on Information Theory*, pages 2591–2595.

Benjamini, Y. and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Ann. Stat.*, 29(4):1165–1188.

Bickel, P. and Levina, E. (2008). Covariance regularization via thresholding. *Ann. Statist.*, 34(6):2577–2604.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning.* Springer Science and Business Media, LLC., Oxford, England.

Browne, M. and Shapiro, A. (1986). The asymptotic covariance matrix of sample correlation coefficients under general conditions. *Linear Algebra and its Applications*, 82:169–176.

Cai, T. T. and Jiang, T. (2011). Limiting laws of coherence of random matrices with applications to testing covariance structure and constructions of compressed sensing materials. *Ann. Statist.*, 39(3):1496–1525.

Cai, T. T., Liu, W., and Xia, Y. (2014). Two-sample covariance matrix testing and support recovery in high-dimensional and sparse settings. *Journ. of the Am. Stat. Ass.*, 108:265–277.

Cai, T. T. and Zhang, A. (2014). Inference on high-dimensional differential correlation matrix. *Technical Report.*

Cai, T. T., Zhang, C.-H., and Zhou, H. H. (2010). Optimal rates of convergence for covariance matrix estimation. *Ann. Statist.*, 38(4):2118–2144.

Choi, Y. and Kendziorski, C. (2009). Statistical methods for gene set co-expression analysis. *Bioinformatics*, 25(21):2780–2786.

Cui, X. and Churchill, G. A. (2003). Statistical tests for differential expression in cdna microarray experiments. *Genome Biology*, 4(4):210.

Datta, S. and Datta, S. (2002). Comparisons and validation of statistical clustering techniques for microarray gene expression data. *Bioinformatics*, 19(4).

Derado, G., Bowman, F. D., and Kilts, C. D. (2010). Modeling the spatial and temporal dependence in fmri data. *Biometrics*, 66(3):949–957.

Donner, A. and Zou, G. (2014). Testing the equality of dependent intraclass correlation coefficients. *J. R. Statist. Soc. D*, 51(3):367–379.

Friguet, C., Kloareg, M., and Causeur, D. (2012). A factor model approach to multiple testing under dependence. *J. Am. Statist. Ass.*, 104(488):1406–1415.

Fukushima, A. (2013). Diffcorr: An r package to analyze and visualize differential correlations in biological networks. *Gene 518*, pages 209–214.

Gill, R., Datta, S., and Datta, S. (2010). A statistical framework for differential network analysis from microarray data. *BMC Bioinformatics*, 11(95):1471–2105.

Greicius, M. D., Krasnow, B., Reiss, A. L., and Menon, V. (2002). Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *P.N.A.S*, 100(1):253–258.

Harman, H. H. (1960). *Modern factor analysis*. Univ. of Chicago Press, New York.

Hu, R., Qiu, X., and Glazko, G. (2010). A new gene selection procedure based on the covariance distance. *Bioinformatics*, 26(3):348–354.

Iglesia, M. D., Vincent, B. G., Parker, J. S., Hoadley, K. A., Carey, L. A., Perou, C. M., and Serody, J. S. (2014). Prognostic b-cell signatures using mrna-seq in patients with subtype-specific breast and ovarian cancer. *Clinical Cancer Research*, 20(14):3818–3829.

Jiang, D., Tang, C., and Zhang, A. (2004). Cluster analysis for gene expression data: A survey. *Knowledge and Data Engineering, IEEE Transactions*, 16.11:1370–1386.

Kriegel, H.-P., Krger, P., and Zimek, A. (2009). Clustering high-dimensional data: A survey on subspace clustering, pattern-based clustering, and correlation clustering. *ACM Trans. on Knowledge Disc. from Data (TKDD)*, 3(1).

Langfelder, P. and Horvath, S. (2008). Wgcna: an r package for weighted correlation network analysis. *BMC bioinformatics*, 9(1):559.

Larntz, K. and Perlman, M. D. (1985). A simple test for the equality of correlation matrices. *Technical Report: University of Washington*, 63:1.

Lewis, K., Kaufman, J., Gonzalez, M., Wimmer, A., and Christakis, N. (2008). Tastes, ties, and time: A new social network dataset using facebook.com. *Social Networks*, 30(4):330–342.

Liu, B.-H., Yu, H., Tu, K., Li, C., Li, Y.-X., and Li, Y.-Y. (2010). Dcgl: an r package for identifying differentially coexpressed genes and links from gene expression microarray data. *Bioinformatics*, 26(20):2637–2638.

MacMahon, M. and Garlaschelli, D. (2015). Community detection for correlation matrices. *Physical Review X*, 5(2).

Muirhead, R. J. (1982). *Aspects of Multivariate Statistical Theory*. John Wiley and Sons, Inc.

Peng, J., Wang, P., Zhou, N., and Zhu, J. (2008). Partial correlation estimation by joint sparse regression models. *J. Am. Statist. Ass.*, 104(486).

Perou, C. M., Srlie, T., Eisen, M. B., van de Rijn, M., Jeffrey, S. S., Rees, C. A., Pollack, J. R., Ross, D. T., Johnsen, H., Akslen, L. A., Fluge, O., Pergamenschikov, A., Williams, C., Zhu, S. X., Lonning, P. E., Borresen-Dale, A.-L., Brown, P. O., and Botstein, D. (2000). Molecular portraits of human breast tumours. *Nature*, 406:747–752.

Phan, K. L., Wager, T., Taylor, S. F., and Liberzon, I. (2002). Functional neuroanatomy of emotion: A meta-analysis of emotion activation studies in pet and fmri. *NeuroImage*, 16:331–348.

Rajaratnam, B., Massam, H., and Carvalho, C. (2008). Flexible covariance estimation in graphical models. *Ann. Statist.*, 36:2818–2849.

Soneson, C. and Delorenzi, M. (2013). A comparison of methods for differential expression analysis of rna-seq data. *BMC Bioinformatics*, 14(91).

Stamatatos, E. (2009). A comparison of methods for differential expression analysis of rna-seq data. *Journ. of the Am. Soc. for Info. Science and Tech.*, 60(3):538–556.

Steiger, J. H. (1980). Tests for comparing elements of a correlation matrix. *Psych. Bulletin*, 87(2):245–251.

Steiger, J. H. and Hakstian, A. R. (1982). The asymptotic distribution of elements of a correlation matrix: Theory and application. *Brit. Journ. of Math. and Stat. Psych.*, 35:208–215.

Voets, N., Adcock, J., Flitney, D., Behrens, T., Hart, Y., Stacey, R., Carpenter, K., and Matthews, P. (2006). Distinct right frontal lobe activation in language processing following left hemisphere injury. *Brain*, 129(3):754–766.

Wainer, H. and Braun, H. I. (2013). *Test Validity*. Routledge.

Wang, J., Fan, L., Wang, Y., Xu, W., Jiang, T., Fox, P. T., Eickhoff, S. B., Yu, C., and Jiang, T. (2015). Determination of the posterior boundary of wernicke's area based on multimodal connectivity profiles. *Human Brain Mapping*, 36:1908–1924.

Wilson, J. D., Wang, S., Mucha, P. J., Bhamidi, S., and Nobel, A. B. (2014). A testing based extraction algorithm for identifying signicant communities in networks. *The Annals of Applied Statistics*, 8(3):1853–1891.

Xia, Y., Cai, T., and Cai, T. T. (2014). Testing differential networks with applications to detecting gene-by-gene interactions. *Biometrika (to appear)*.

Zhou, C., Han, F., Zhang, X., and Liu, H. (2015). An extreme-value approach for testing the equality of large u-statistic based correlation matrices. *arXiv:1502.03211*.